

A Bi-Illuminant Dichromatic Reflection Model for Understanding Images

Bruce A. Maxwell
Tandent Vision Science, Inc.
brucemaxwell@tandent.com

Richard M. Friedhoff
Tandent Vision Science, Inc.
richardfriedhoff@tandent.com

Casey A. Smith
Tandent Vision Science, Inc.
caseysmith@tandent.com

Abstract

This paper presents a new model for understanding the appearance of objects that exhibit both body and surface reflection under realistic illumination. Specifically, the model represents the appearance of surfaces that interact with a dominant illuminant and a non-negligible ambient illuminant that may have different spectral power distributions. Real illumination environments usually have an ambient illuminant, and the current dynamic range of consumer cameras is sufficient to capture significant information in shadows. The bi-illuminant dichromatic reflection model explains numerous empirical findings in the literature and has implications for commonly used chromaticity spaces that claim to be illumination invariant but are not in many natural situations. One outcome of the model is the first 2-D chromaticity space for an RGB image that is robust to illumination change given dominant and ambient illuminants with different spectral power distributions.

1. Introduction

The light emitted from an object is a complex combination of its material properties and the illumination environment. A significant amount of prior work has focused on describing material properties, but work on physics-based appearance has generally assumed a single dominant illuminant in the scene and either ignored ambient illumination or considered it insignificant [13, 12, 14, 11, 17, 19]. Conversely, work on understanding interreflection or diffuse illumination has often used simple models of material properties [20, 16].

Understanding the appearance of real surfaces is critical to tasks such as segmentation, object recognition, and material recognition. Most approaches to understanding color appearance assume that the shape of an object's color spectrum is invariant to illumination change; as the intensity of the light changes, the intensity of the object changes but its chromaticity, or the relative power in each color band, remains constant. Unfortunately, the assumption is false in most real-world scenes.



Figure 1. (a) Original image taken outside, (b) chromaticity. Objects change color in shadows, complicating scene understanding.

Consider figure 1, a scene illuminated by typical outdoor illuminants: sunlight and skylight. The light in the fully lit areas is dominated by the sun while the light in the shadows is dominated by blue sky. Therefore, as the illumination intensity decreases, the measured color of the surfaces becomes more blue. Because of the linkage between illumination level and color, segmenting the image using traditional measures such as intensity, RGB, chromaticity, or hue-saturation will separate the flower shadows on the concrete from the lit concrete. The connection between illumination level and color is not random, however, but instead is a function of the illumination environment and material properties. This work was inspired by and builds upon experiments by Friedhoff at the Rowand Institute for Science in 1994 showing the spectral ratio of lit and shadowed areas is consistent across materials in a planar scene. By modeling and analyzing realistic illumination environments and material properties we can explain a number of observed phenomena and develop new tools for understanding the appearance of objects in real environments.

2. Related Work

The goal of understanding, analyzing, and differentiating material and illumination properties has a long history in computer vision. Land and McCann's retinex was one of the first algorithms proposed for separating material and illumination effects, followed by Horn, who showed how to implement 2D removal of small gradients [15][13]. Barrow and Tenenbaum argued more generally that separating an image into intrinsic parts—i.e. reflectance, illumination,

range, and orientation—would aid in image understanding [3]. Color spaces derived from physical models that help to differentiate the different intrinsic attributes are an essential part of this task.

Shafer’s dichromatic reflection model was the first physics-based model of reflection to separate different types of reflection on a single surface [22]. Shafer’s model proposed that inhomogeneous dielectrics such as paints, ceramics and plastics exhibit two types of reflection, surface reflection and body reflection, and that the different types of reflection caused specific types of changes in appearance. Klinker et al. used the model to develop an algorithm for separating surface and body reflection in controlled conditions with a single illuminant [14].

While Shafer’s original paper on the dichromatic reflection model included a term for ambient illumination, he did not separate its effects into surface and body reflection, and Klinker et al. removed the term from the model for their experiments because in the controlled conditions the ambient illumination was effectively zero [14].

Gershon et al. proposed a similar model, but divided the reflection into three components: specular, ambient, and diffuse [10]. These reflection modes interacted with direct and ambient illuminants, but not in an orthogonal manner. Where Shafer’s material reflection model was based on a physical analysis of inhomogeneous dielectrics, Gershon’s model arbitrarily divides the reflection components, linking the type of reflection with the type of illumination, despite the fact that the mechanisms of reflection work similarly for all visible light energy, regardless of its source.

Healey proposed a unichromatic reflection model for metals and was able to extract useful information about the illumination, reflection and material type under controlled conditions and a single illuminant [12]. Even more recent work by Mallick et. al. and Narasimhan et. al. on physics-based color spaces still assumed a single illuminant for the scene [17, 19].

Barnard et al. were some of the first to explicitly consider two dominant illuminants [2]. They presented an algorithm for color constancy under two illuminants of approximately equal intensity but different colors.

In summary, except for [22] and [10], the early work in understanding appearance assumed a single light source in the scene. Shadows were disregarded, even by Shafer, in part because the dynamic range and noise levels of cameras limited the utility of color measurements in shadows [22]. Thus, while basic computer graphics models have included an ambient term for decades, in computer vision the ambient term has been all but ignored for analysis.

More recent work has observed some of the impacts of an ambient illuminant on appearance. Funke-Lea et al. used a model mathematically similar to Shafer’s extended dichromatic model and noted that the body reflection

of a surface with varying illumination could be represented as a line segment [9].

The color lines approach of Omer observes that curves in RGB space that correspond to a single object’s appearance are offset from the origin and do not necessarily intersect the origin if extended [21]. Shafer predicted both of these observations in the first dichromatic reflection model paper.

Weijer and Schmid resurrected Shafer’s extended dichromatic reflection model to generate color features for matching that provide illumination invariance [25]. While most of their resulting features focused on only the body reflection terms under a single illuminant, they do propose one color invariant that works for the more general case of a direct and an ambient illuminant.

Barnard and Finlayson proposed a method for identifying shadows using the simple ratio of pixels on a material lit only by ambient illumination to pixels on the same material lit by both ambient and direct illumination [1]. While a proof is not provided in the paper, they claimed that the ratios would be similar across materials and used that premise, along with other constraints, to differentiate potential shadow boundaries from material boundaries.

In this paper we fully extend Shafer’s dichromatic reflection model to handle an ambient and a direct illuminant and show that the phenomena observed by Omer and Barnard and Finlayson are predicted and understandable using the new model. We also demonstrate the difficulty of creating a true chromaticity space given a two-illuminant situation. Most chromaticity spaces, including traditional ones like hue-saturation and rg-chromaticity, work only if there is a single illuminant, as shown in section 4.

Work by Marchant and Onyango and an equivalent method by Finlayson et al. presented a 1-D chromaticity space for RGB that is invariant to any single Planckian illuminant [18, 7, 6]. Using our model, we derive a 2-D chromaticity space that is robust to any bi-illuminant situation (direct + ambient), including non-Planckian illuminants.

In summary, the bi-illuminant dichromatic reflection model we propose extends Shafer’s model of appearance, explains phenomena observed in related work, explains the failure of traditional chromaticity spaces in natural imagery, and provides a comprehensive framework for analyzing reflection from common materials in imagery taken under realistic lighting conditions.

3. Bi-Illuminant Dichromatic Reflection Model

To derive our bi-illuminant model, we first review Shafer’s dichromatic reflection model, which separates reflection from inhomogeneous dielectric surfaces into two parts: body reflection and surface reflection. Body reflection occurs when light interacts with pigment particles within a material, modifying the spectral distribution of the

light and reflecting it in random directions. Surface reflection occurs at the interface between two materials of different density and is generally highly directional. The two types of reflection have different effects on the color and intensity of an object's appearance. The validity of the dichromatic reflection model has been extensively tested [23, 24].

Shafer's extended dichromatic reflection model expresses the radiance of a surface as three terms: the body reflection, the surface reflection, and an ambient term that is constant for each material [22]. Using this model, and assuming a given geometry between the light source and the surface, the radiance of a surface point in the direction given by the pan and tilt angles (θ_e, ϕ_e) can be expressed as in (1), where m_b and c_b are the achromatic and chromatic aspects of the body reflection, m_s and c_s are the achromatic and chromatic aspects of the surface reflection and c_d is the contribution of diffuse illumination.

$$I(\theta_e, \phi_e, \lambda) = m_b(\theta_e, \phi_e)c_b(\lambda) + m_s(\theta_e, \phi_e)c_s(\lambda) + c_d(\lambda) \quad (1)$$

The dichromatic reflection model says that both the body and surface reflection components can be divided into achromatic and chromatic parts. The achromatic terms m_b and m_s express the magnitude of the body reflection and surface reflection, respectively. The chromatic terms c_b and c_s express how the color of the incident light is modified by the two different types of reflection.

In Shafer's extended dichromatic reflection model, the diffuse illumination is assumed to be constant across a surface and over the illumination hemisphere, so c_d is constant for a given surface. Shafer did not expand the diffuse term and assumed it would be negligible given the limitations on the dynamic range and noise levels of sensors when the paper was published. In almost all later work based on Shafer's model, the diffuse term was left out and experiments assumed a single illuminant in the scene [14][12].

The extended dichromatic reflection model is incomplete, however. Locally within most scenes there are two illuminants. One illuminant, which we define as the direct illuminant, is the primary illuminant in the scene, and it causes perceptible shadows if blocked. The ambient illuminant is the secondary illuminant and is the result of the reflection of the direct illuminant from all of the reflecting or scattering surfaces in the scene. A common bi-illuminant situation is an outdoor scene with a yellow sun and a blue sky, but, as noted by Shafer, almost all typical indoor scenes also exhibit two illuminant behavior [22].

Both the direct and ambient illuminants produce light that reflects from a surface, and surfaces exhibit surface and body reflection regardless of the illuminant. Explicitly modeling both ambient surface reflection and ambient body reflections lets us analyze their effects independently.

Current generation digital cameras have sufficiently high dynamic range (12 bits/pixel) and low noise that the effects of ambient body and surface reflection are measurable. Ambient reflection characteristics will only increase in importance as dynamic range, resolution, and signal-to-noise ratios improve. Furthermore, their effects violate existing methods of measuring color similarity, causing traditional models of chromaticity, for example, to fail in real images.

A true bi-illuminant dichromatic reflection model is necessary to understand and analyze colors in real imagery. It allows us to take advantage of new sensor capabilities and develop new ways of thinking about color similarity based on realistic physical models of appearance.

We propose a new bi-illuminant dichromatic reflection [BIDR] model that models the two illuminants orthogonally from the dichromatic reflection function of the surface. The latter is a bi-directional reflectance distribution function [BRDF], which we can divide into body and surface reflection components, each of which contains an achromatic and chromatic part, as in the dichromatic reflection model. The BRDF describes the amount of illumination reflected in a particular angle or emittance (θ_e, ϕ_e) given illumination at a particular angle of incidence (θ_i, ϕ_i) . In the case of a single primary illuminant, the incident angle is the direction of the primary illuminant relative to the surface.

Conceptually, the BIDR model requires four terms: two types of reflection for the direct illuminant and two types of reflection for the ambient illumination. The types of illumination act differently, however, as the direct illuminant is normally a small area source, while the ambient is integrated over the entire hemisphere. For this derivation we assume that the chromatic term of the BRDF for ambient reflection is isotropic. In other words, if the material modifies the spectrum of the incident light, that modification is the same for all incident and viewing angles for the ambient illuminant. While this is not the case for materials such as copper or bronze, the effect on the ambient term color is likely to be small. Given those assumptions, the complete BIDR model is given in (2), with m_b and m_s representing the achromatic reflection characteristics and c_b and c_s representing the chromatic reflection characteristics as in (1).

$$I(\theta_e, \phi_e, \lambda) = m_b(\theta_e, \phi_e, \theta_i, \phi_i)c_b(\lambda)l_d(\theta_L, \phi_L, \lambda) + m_s(\theta_e, \phi_e, \theta_i, \phi_i)c_s(\lambda)l_d(\theta_L, \phi_L, \lambda) + c_b(\lambda) \int_{\theta_i, \phi_i} m_b(\theta_e, \phi_e, \theta_i, \phi_i)l_a(\theta_i, \phi_i, \lambda)d\theta_i, d\phi_i + c_s(\lambda) \int_{\theta_i, \phi_i} m_s(\theta_e, \phi_e, \theta_i, \phi_i)l_a(\theta_i, \phi_i, \lambda)d\theta_i, d\phi_i \quad (2)$$

The direct illuminant l_d in direction (θ_L, ϕ_L) relative to the local surface normal multiplies the dichromatic BRDF

to produce two terms corresponding to the traditional body and surface reflection components. The range of m_b , m_s , c_b , and c_s in this derivation is $[0, 1]$, as they each represent the fraction of incoming light reflected by the material. The ambient illuminant l_a must be integrated over the visible hemisphere on the surface, except in the direction of the direct illuminant. The assumption of isotropic chromatic BRDF terms for the ambient illuminant reflection means that $c_b(\lambda)$ and $c_s(\lambda)$ can be factored out of the body and surface reflection integrals.

If we represent the two ambient integrals as $M_{ab}(\theta_e, \phi_e, \lambda)$ and $M_{as}(\theta_e, \phi_e, \lambda)$, respectively, then the BIDR model is given as (3).

$$I(\theta_e, \phi_e, \lambda) = m_b(\theta_e, \phi_e, \theta_i, \phi_i)c_b(\lambda)l_d(\theta_L, \phi_L, \lambda) + m_s(\theta_e, \phi_e, \theta_i, \phi_i)c_s(\lambda)l_d(\theta_L, \phi_L, \lambda) + c_b(\lambda)M_{ab}(\theta_e, \phi_e, \lambda) + c_s(\lambda)M_{as}(\theta_e, \phi_e, \lambda) \quad (3)$$

Given a specific geometry, the scene angles become constants, which simplifies the BIDR model to (4).

$$I(\lambda) = m_b c_b(\lambda) l_d(\lambda) + m_s c_s(\lambda) l_d(\lambda) + c_b(\lambda) M_{ab}(\lambda) + c_s(\lambda) M_{as}(\lambda) \quad (4)$$

If we are interested in the appearance of cast shadows on a surface, we can add a shadow term $H \in [0, 1]$ that multiplies the direct illuminant term and indicates the percent blocked, allowing the formation of umbra and penumbra. When the direct illuminant is completely blocked, $H = 0$ and the only surface radiance comes from reflection of the ambient illumination.

$$I(\lambda) = [m_b c_b(\lambda) l_d(\lambda) + m_s c_s(\lambda) l_d(\lambda)] H + c_b(\lambda) M_{ab}(\lambda) + c_s(\lambda) M_{as}(\lambda) \quad (5)$$

As the shadow term is achromatic, we can combine it with the geometric terms to get the achromatic multipliers $\gamma_b = m_b H$ and $\gamma_s = m_s H$. Substituting, we get (6).

$$I(\lambda) = \gamma_b c_b(\lambda) l_d(\lambda) + \gamma_s c_s(\lambda) l_d(\lambda) + c_b(\lambda) M_{ab}(\lambda) + c_s(\lambda) M_{as}(\lambda) \quad (6)$$

For a matte surface, the surface reflection terms are negligible, giving the simplified form of the BIDR model in (7) that includes only the body reflection terms.

$$I(\lambda) = \gamma_b c_b(\lambda) l_d(\lambda) + c_b(\lambda) M_{ab}(\lambda) \quad (7)$$

The BIDR model as defined provides a description of the radiance leaving a surface. In order to convert the surface

radiance into an image, the signal must go through a sensor. The output of a sensor C_i with a response curve $F_i(\lambda)$ and input radiance signal $I(\lambda)$ is given by (8), integrated over the visible wavelengths.

$$C_i = \int F_i(\lambda) I(\lambda) d\lambda \quad (8)$$

Substituting (6) into (8), we get (9).

$$C_i = \gamma_b \int F_i(\lambda) c_b(\lambda) l_d(\lambda) d\lambda + \gamma_s \int F_i(\lambda) c_s(\lambda) l_d(\lambda) d\lambda + \int F_i(\lambda) c_b(\lambda) M_{ab}(\lambda) d\lambda + \int F_i(\lambda) c_s(\lambda) M_{as}(\lambda) d\lambda \quad (9)$$

In a digital camera, there are typically three sensor values per pixel, so $i \in R, G, B$, but the derivation applies equally well to images with more sensor bands.

An approximation to (9) can be obtained by assuming the sensors are sharp enough (close to delta functions), or the spectrum of the radiance is smooth enough (close to constant across the sensor response curve), to treat the integral as a multiplication, which is not unreasonable [4]. The resulting expression in (10) provides a complete description of the value of a pixel in an image corresponding to a location in the scene exhibiting both body and surface reflection from both the direct and ambient illuminant. The superscripts for c_b , c_s , l_d , M_{ab} , and M_{as} indicate the terms represent the part of the spectrum that intersects with F_i .

$$C_i = \gamma_b F_i c_b^i l_d^i + \gamma_s F_i c_s^i l_d^i + F_i c_b^i M_{ab}^i + F_i c_s^i M_{as}^i \quad (10)$$

The equivalent expression for only the body reflection model is given in (11).

$$C_i = \gamma_b F_i c_b^i l_d^i + F_i c_b^i M_{ab}^i \quad (11)$$

The expression in (11) defines a line segment in color space (e.g. RGB) with a fully shadowed ($\gamma_b = 0$) pixel at one end, a fully lit ($\gamma_b = 1$) pixel at the other end and penumbra in the middle. Given noise in the material color and the imaging system, in practice a material under varying illumination defines a cylinder, which we denote as bi-illuminant dichromatic reflection [BIDR] cylinders.

There are three key observations that result from (11). First, the infinite extension of the BIDR cylinder ($\gamma \in [-\infty, \infty]$) does not intersect the origin unless the normalized spectra of the ambient and direct illuminants are identical. Second, because the ambient reflection is material dependent, there is no single offset that will cause all material cylinders to intersect the origin. Third, BIDR cylinders representing two different materials with differing amounts of direct illumination can intersect without being collinear.

For example, a material in 20% direct illumination can be identical to a different material in 80% direct illumination.

One implication of the above three observations is that traditional measures of color such as hue-saturation and normalized color are not invariant to changes in illumination intensity, a fact that has been noted by others [8]. A surface with a uniform body reflection will change color as the intensity of the direct illuminant changes. As shown in figure 5, a surface of constant material reflection can change its hue from red to blue as the amount of direct illuminant varies under natural illumination conditions.

3.1. Spectral Ratio

Given the BIDR model, it is possible to identify measures of scene properties that are invariant even under two illuminants with significantly different power spectra. In particular, consider the BIDR cylinder defined by (11). Each material has its own BIDR cylinder. The dark end is the product of the body reflection and the ambient illuminant, and the orientation of the cylinder is the product of the body reflection and the direct illuminant.

Given that the material's body reflection is common to both terms, the ratio of the cylinder's dark end to the cylinder's orientation measures the ratio of illuminants. Therefore, we can calculate a ratio of observed lit and shadowed pixels located on the same material and obtain a ratio of illuminants. We define the ratio of illuminants to be the spectral ratio \vec{S} , where for an RGB image $\vec{S} = (S_R, S_G, S_B)$.

$$\begin{aligned} S_i &= \frac{\text{dark end}}{\text{orientation}} = \frac{\text{shadowed}^i}{\text{lit}^i - \text{shadowed}^i} \\ &= \frac{F_i c_b^i M_{ab}^i}{(F_i c_b^i l_d^i + F_i c_b^i M_{ab}^i) - F_i c_b^i M_{ab}^i} \quad (12) \\ &= \frac{M_{ab}^i}{l_d^i} = \frac{\text{ambient}^i}{\text{direct}^i} \end{aligned}$$

Given several different materials on a surface, such as on a picture book, if a shadow were cast across the book we would expect the dark and lit pixels of each material to exhibit similar spectral ratios. Figure 3 demonstrates an example of similar spectral ratios on a single surface.

One problem with measuring raw spectral ratios across a surface is that the γ_b term for a lit pixel can vary significantly due to changes in either geometry or partial shadowing. A raw spectral ratio will be consistent for two different materials only if the bright pixels on each material have identical γ_b values.

A normalized spectral ratio, however, exhibits invariance to the γ_b value, so long as there is sufficient difference between the bright and dark pixels to robustly measure the orientation of the cylinder. (13) gives the spectral ratio normalized by its length D . As shown in (14) for RGB data, the

γ_b term can be extracted from the normalization constant, and if we substitute (14) into (13), the γ_b terms cancel.

$$\hat{S}_i = \frac{\text{shadowed}}{\text{lit} - \text{shadowed}} \left(\frac{1}{D} \right) = \frac{M_{ab}^i}{\gamma_b l_d^i} \left(\frac{1}{D} \right) \quad (13)$$

$$\begin{aligned} D &= \sqrt{\left(\frac{M_{ab}^R}{\gamma_b l_d^R} \right)^2 + \left(\frac{M_{ab}^G}{\gamma_b l_d^G} \right)^2 + \left(\frac{M_{ab}^B}{\gamma_b l_d^B} \right)^2} \\ &= \frac{1}{\gamma_b} \sqrt{\left(\frac{M_{ab}^R}{l_d^R} \right)^2 + \left(\frac{M_{ab}^G}{l_d^G} \right)^2 + \left(\frac{M_{ab}^B}{l_d^B} \right)^2} \quad (14) \end{aligned}$$

Therefore, we would expect the normalized spectral ratio to be constant across materials under three conditions. First, the ambient and direct illuminants must be constant across the materials. Second, the dark pixel used in the expression must represent a $\gamma_b = 0$ pixel. Third, there must be sufficient difference between the dark and bright pixels to robustly measure the orientation of the BIDR cylinder.

3.2. Log Space Chromaticity

The appearance of a surface as the direct illumination varies in brightness forms a BIDR cylinder in RGB space. The cylinder for each material points in a unique direction because the vector of change is the product of the material color and direct illuminant.

We can rewrite the BIDR cylinder equation (11) as in (15), which describes the appearance of the surface as the product of its illumination and material properties. Note that the only term that changes as the amount of direct illumination changes is the value of γ_b .

$$C_i = F_i c_b^i (\gamma_b l_d^i + M_{ab}^i) = F_i c_b^i M_{ab}^i \left(\frac{\gamma_b l_d^i}{M_{ab}^i} + 1 \right) \quad (15)$$

Taking the log of (15), we get (16), where S_i is the spectral ratio as defined in (12).

$$\begin{aligned} \log C_i &= \log \left(F_i c_b^i M_{ab}^i \left(\frac{\gamma_b l_d^i}{M_{ab}^i} + 1 \right) \right) \\ &= \log F_i c_b^i + \log M_{ab}^i + \log \left(\frac{\gamma_b}{S_i} + 1 \right) \quad (16) \end{aligned}$$

The important thing to note is that only the third term in (16) varies with change in the amount of direct illumination falling on the surface. The first term is a material dependent constant and the second term is an ambient illumination constant. Therefore, the shape of the curve representing a BIDR cylinder in log space is completely determined by the ratio of direct and ambient illuminants regardless of the

material color. Different materials produce curves that are translated versions of one another.

Furthermore, for spectral ratios that are not strongly saturated—where the spectra of the direct and ambient illuminants are not exceptionally different—the curve segment in log space is close to linear, as shown in section 4.

The fact that the curves are translated (and therefore parallel) versions of one another and close to linear suggests a new chromaticity space: the plane perpendicular to the lines connecting the dark and bright points of the log space curves. In log space, the orientation defined by $\vec{N} = \log(\text{bright}) - \log(\text{dark})$, is invariant to material color.

$$\begin{aligned}
 N_i &= \log(\text{bright}) - \log(\text{dark}) \\
 &= \left[\log F_i c_b^i + \log M_{ab}^i + \log \left(\frac{1}{S_i} + 1 \right) \right] - \\
 &\quad \left[\log F_i c_b^i + \log M_{ab}^i + \log \left(\frac{0}{S_i} + 1 \right) \right] \quad (17) \\
 &= \log \left(\frac{1}{S_i} + 1 \right)
 \end{aligned}$$

Given a plane perpendicular to a linear approximation of the curves, the log space chromaticity is defined as the projection of the log of the sensor values onto the plane defined by \vec{N} as in (17). Note that a zero sensor reading in any channel means the log space chromaticity is undefined.

If the spectral ratio of a scene is known, then the log space chromaticity plane orientation is defined, and every pixel can be converted into a 2D chromaticity.

If the spectral ratio of a scene is unknown, the log space chromaticity plane can still be estimated. One approach is to use an entropy minimization method to identify the proper plane orientation, similar to Finlayson’s approach to identifying a 1-D illumination invariant space, but with the search in two dimensions [6]. The entropy of the projection of log pixel values onto the plane will, in general, be minimized when the log space curves align, projecting all of the pixels on a single material into a small cluster.

Note that there are two benefits to our log space chromaticity when compared to prior work. First, our chromaticity space is explicitly built to handle two illuminants, even when those illuminants are not Planckian. Second our chromaticity space is a true 2D chromaticity, as opposed to the 1D chromaticity space of Marchant and Onyango or Finlayson et al. [18, 7]. While Drew et al. reproject the 1D chromaticity back to 2D by estimating what cluster of colors each pixel came from, their chromaticity data is inherently one dimensional [5].

4. Experiments

All images used in the following experiments were captured using RAW mode on a Canon 20D under natural light-




Figure 2. (a) Original image taken outside, (b) histogram in RGB space of the two selected areas crossing from fully lit to shadow. Single materials have colors which form straight lines in RGB space. Here, these lines do not intersect the origin.

ing conditions, mostly daylight. After capture, the images were converted into 16-bit tiffs to preserve the original dynamic range of the sensor (≈ 12 bits/pixel) and the linearity of the sensor values.

4.1. Properties of the BIDR model

The first experiment just looks at the shape of the curves in RGB space for several different materials. The RGB histogram in Figure 2(b) of the extracted sections of image in Figure 2(a) demonstrates two important properties of BIDR cylinders. The cylinders are all offset from the origin by a different amount and their infinite extensions do not intersect the origin.

Our second experiment looks at the constancy of spectral ratios across multiple materials. The two images in figures 3(a) and (b) are the same scene taken outdoors on a sunny day with a clear sky. The two images in figures 3(c) and (d) were taken indoors with reddish incandescent lights providing the ambient illumination and a neutral direct illuminant. In both cases the images were captured a few seconds apart, but with a cast shadow on the object in the second image. After aligning the images, we masked out pixels that were similar (no shadow) and pixels with strong gradients in the lit image, as they were likely to be nonuniform blends of materials and the image registration was not perfect. Finally, we calculated the spectral ratio between the shadowed image and the unshadowed image for all the unmasked pixels. The resulting spectral ratios are shown visually in figures 3(e) and (f). The mean normalized spectral ratio for unmasked pixels in the outdoor image was $\hat{S} = (\hat{S}_R, \hat{S}_G, \hat{S}_B) = (0.49, 0.57, 0.66)$ with standard deviations of $\vec{\sigma} = (0.023, 0.017, 0.022)$. The mean normalized spectral ratio for the indoor scene was $\hat{S} = (0.80, 0.53, 0.27)$ with standard deviations of $\vec{\sigma} = (0.018, 0.018, 0.024)$. Both scenes demonstrate consistent normalized spectral ratios across materials with different colors, despite the differences in illuminants.

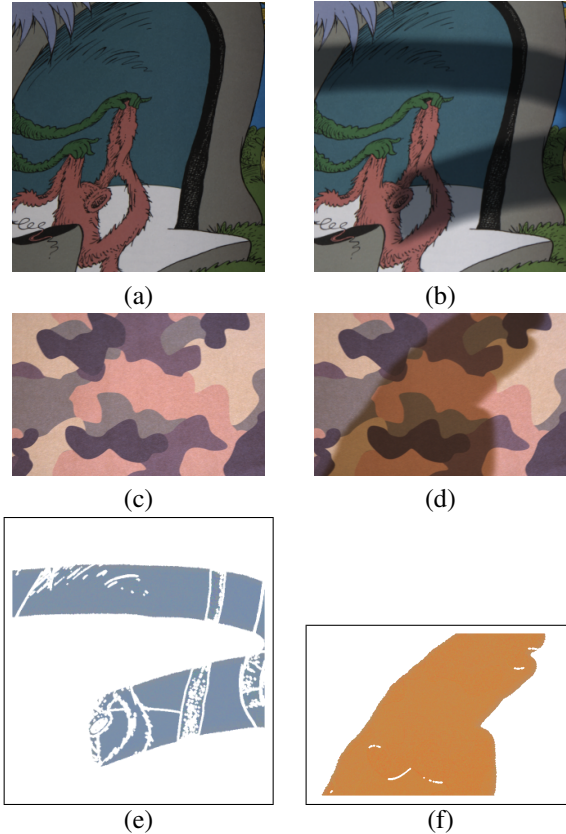


Figure 3. (a) & (b) Scene 1 with and without a shadow (c) & (d) scene 2 with and without a shadow, (e) & (f) normalized spectral ratios calculated for shadowed pixels in uniform areas. In each scene, the spectral ratio is constant regardless of the material or illuminant colors.

4.2. Log space chromaticity

Figures 4(a) and (c) show an image taken outdoors and its log space histogram. Note the almost linear and parallel nature of the lines corresponding to each material. The normalized spectral ratio for this scene is mildly blue $\hat{S} = (0.49, 0.57, 0.66)$. In contrast, the outdoor image in figure 4(b) has a very saturated normalized spectral ratio of $\hat{S} = (0.28, 0.43, 0.86)$, and the log space histogram in figure 4(d) exhibits slight curvature.

Figures 5 and 6 show two scenes in RGB, hue-saturation, normalized color, and our log space chromaticity. For each scene, the log space chromaticity plane was found by an exhaustive 2-D search on valid plane orientations to minimize the entropy of the projection of scene colors onto the plane.

Figure 7 shows the flower image in our log space chromaticity with the spectral ratio selected for the outer edge of the shadows. The interreflection from the flowers causes the ambient illumination to change underneath them, making the shadows there more visible. The scene demonstrates local, but not global consistency of spectral ratios.




Figure 4. (a) Original image taken outside, (b) original image from late afternoon (c) histogram of (a) in $\log(RGB)$ space, (d) histogram of (b) in $\log(RGB)$. Measured colors of different materials form parallel curves in $\log(RGB)$ space.




Figure 5. (a) Original image taken outside, (b) hue-saturation, (c) normalized color, (d) log space chromaticity. Note that while the shadow is boldly visible in hue-saturation and normalized color, it is nearly invisible in log space chromaticity.

5. Summary

Understanding images requires that we understand how materials appear under realistic illumination conditions. Absent that understanding, tasks such as segmentation and object recognition become more difficult because simple assumptions about how material colors behave under varying illumination create apparently random effects.

The bi-illuminant dichromatic reflection model describes the appearance of materials interacting with a direct light source and an ambient illumination environment via body and surface reflection. The model explains why standard chromaticity methods are not invariant to typical illumination change and is the basis for a new 2-D chromaticity space that is illumination invariant for a direct and ambient illuminant pair with differing spectra. The model is also the basis for a spectral ratio calculation that measures the properties of the light field irrespective of the underlying material reflectances in the scene.

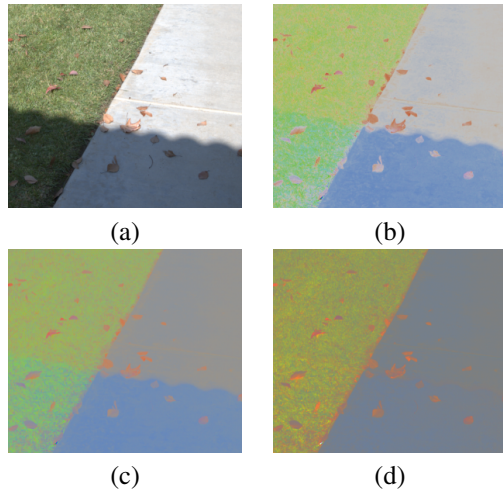


Figure 6. (a) Original image taken outside, (b) hue-saturation, (c) normalized color, (d) log space chromaticity. Note that while the shadow is boldly visible in hue-saturation and normalized color, it is nearly invisible in log space chromaticity.



Figure 7. Log space chromaticity image of the flower scene from figure 1. Note the shadows are visible where interreflection from the flowers changes the ambient illumination.

References

- [1] K. Barnard and G. D. Finlayson. Shadow identification using colour ratios. In *The 8th Color Imaging Conf.*, Scottsdale, AZ, 2000.
- [2] K. Barnard, G. D. Finlayson, and B. Funt. Color constancy for scenes with varying illumination. *Computer Vision and Image Understanding*, 65(2):311–321, February 1997.
- [3] H. G. Barrow and J. M. Tenenbaum. *Computer Vision Systems*, chapter Recovering intrinsic scene characteristics from images, pages 2–26. Academic Press, 1978.
- [4] C. F. Borges. A trichromatic approximation method for surface illumination. *J. of Optical Society of America*, 8(8):1319–1323, August 1991.
- [5] M. S. Drew, G. D. Finlayson, and S. D. Hordley. Recovery of a chromaticity image free from shadows via illumination invariance. In *IEEE Workshop on Color and Photometric Methods in Computer Vision*, Nice, France, 2003.
- [6] G. D. Finlayson, M. S. Drew, and L. Cheng. Intrinsic images by entropy minimization. In T. Pajdla and J. Matas, editors, *Proc. of European Conf. on Computer Vision*, LNCS 3023, pages 582–595, 2004.
- [7] G. D. Finlayson, S. D. Hordley, and M. S. Drew. Removing shadows from images. In *Proc. of European Conf. on Computer Vision*, pages 823–836, London, UK, 2002. Springer-Verlag.
- [8] G. D. Finlayson, S. D. Hordley, C. Lu, and M. S. Drew. On the removal of shadows from images. *IEEE Trans. on Pattern Analysis and Machine Vision*, 28(1):59–68, 2006.
- [9] G. D. Funka-Lea and R. Bajcsy. Combining color and geometry for the active, visual recognition of shadows. In *Proc. Fifth Int'l Conf. on Computer Vision*, Cambridge, 1995.
- [10] R. Gershon, A. D. Jepson, and J. K. Tsotsos. Ambient illumination and the determination of material changes. *J. of the Optical Society of America A*, 3(10):1700–1707, 1986.
- [11] J.-M. Geusebroek, R. v. d. Boomgaard, A. W. M. Smeulders, and H. Geerts. Color invariance. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(12):1338–1350, December 2001.
- [12] G. E. Healey. Using color for geometry-insensitive segmentation. *J. of the Optical Society of America A*, 6(6):920–937, June 1989.
- [13] B. K. P. Horn. Determining lightness from an image. *Computer Graphics and Image Processing*, 3(1):277–299, December 1974.
- [14] G. J. Klunker, S. A. Shafer, and T. Kanade. A physical approach to image understanding. *Int'l J. of Computer Vision*, 4(1):7–38, January 1990.
- [15] E. H. Land and J. J. McCann. Lightness and the retinex theory. *J. of the Optical Society of America*, 61:1–11, 1971.
- [16] M. S. Langer. When shadows become interreflections. *Int'l J. of Computer Vision*, 34(2/3):193–204, 1999.
- [17] S. P. Mallick, T. Zickler, D. J. Kriegman, , and P. N. Belhumeur. Beyond lambert: Reconstructing specular surfaces using color. In *IEEE Proc. of Computer Vision and Pattern Recognition*, pages 619–626, June 2005.
- [18] J. A. Marchant and C. M. Onyango. Shadow-invariant classification for scenes illuminated by daylight. *J. of the Optical Society of America A*, 17(11), November 2000.
- [19] S. G. Narasimhan, V. Ramesh, and S. K. Nayar. A class of photometric invariants: Separating material from shape and illumination. In *Proc. of Int'l Conf. on Computer Vision*, pages 1387–1394, 2003.
- [20] S. Nayar, K. Ikeuchi, and T. Kanade. Shape from Interreflections. In *IEEE Int'l Conf. on Computer Vision (ICCV)*, pages 2–11, Dec 1990.
- [21] I. Omer and M. Werman. Color lines: Image specific color representation. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 946–953, June 2004.
- [22] S. A. Shafer. Using color to separate reflection components. *Color Research Applications*, 10:210–218, 1985.
- [23] S. Tominaga. Surface identification using the dichromatic reflection model. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):658–670, 1991.
- [24] S. Tominaga and N. Tanaka. Estimating reflection parameters from a single color image. *IEEE Comput. Graph. Appl.*, 20(5):58–66, 2000.
- [25] J. van de Weijer and C. Schmid. Coloring local feature extraction. In *European Conf. on Computer Vision*, volume 3952 of LNCS, pages 334–348. Springer, 2006.